14 AI-M-392

UNCLASSIFIED

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AIM-392	**2. GOVT ACCESSION NO.** N00014-75-C0643	**3. RECIPIENT'S CATALOG NUMBER**
4. TITLE (and Subtitle) COMPUTATION OF LOCALLY PARALLEL STRUCTURE.		**5. TYPE OF REPORT & PERIOD COVERED** Memorandum rept.
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) Kent A. Stevens		**8. CONTRACT OR GRANT NUMBER(s)** N00014-75-C-0643
9. PERFORMING ORGANIZATION NAME AND ADDRESS Artificial Intelligence Laboratory 545 Technology Square Cambridge, Massachusetts 02139		**10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS**
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Blvd Arlington, Virginia 22209		**12. REPORT DATE** March 1977
		13. NUMBER OF PAGES 31 33
14. MONITORING AGENCY NAME & ADDRESS(If different from Controlling Office) Office of Naval Research Information Systems Arlington, Virginia 22217		**15. SECURITY CLASS. (of this report)** UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Distribution of this document is unlimited.

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

DDC
RECEIVED
APR 10 1978
F

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

Parallelism
Vision
Perception
Grouping

Dot Patterns
Segmentation
Image Processing

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

A Moire-like effect can be observed in dot patterns consisting of two super-imposed copies of a random dot pattern where one copy has been expanded, translated, or rotated. One perceives in these patterns a structure that is shown by experiment to be limited by the local geometry of the pattern, in-dependent of the overall structure or the dot density. A simple representa-tion of locally parallel structure is proposed, and it is found to be com-putable by a non-iterative, parallel algorithm. An implementation of this

407 483   JOB

CONTINUATION OF BLOCK 20

algorithm is demonstrated. Its performance parallels that observed experimentally, providing a potential explanation for human performance. Advantages are discussed for the early description of locally parallel structure in the course of visual processing.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 392

March 1977

COMPUTATION OF LOCALLY PARALLEL STRUCTURE

by

Kent A. Stevens

ABSTRACT: A Moiré-like effect can be observed in dot patterns consisting of two superimposed copies of a random dot pattern where one copy has been expanded, translated, or rotated. One perceives in these patterns a structure that is locally parallel. Our ability to perceive this structure is shown by experiment to be limited by the local geometry of the pattern, independent of the overall structure or the dot density. A simple representation of locally parallel structure is proposed, and it is found to be computable by a non-iterative, parallel algorithm. An implementation of this algorithm is demonstrated. Its performance parallels that observed experimentally, providing a potential explanation for human performance. Advantages are discussed for the early description of locally parallel structure in the course of visual processing.

Summary

1. A Moiré-like effect can be seen when two copies of a dot pattern are superimposed, where one copy is slightly transformed, e.g., by rotation, or translation. These *Glass patterns* have been taken as evidence that the visual system performs local autocorrelations. It has also been understood that the effect is due to the detection of pairs of correlated dots, and that the pairings need not be nearest neighbors. It had been observed that the concentric structure seen in rotation-generated patterns vanishes as the degree of rotation increases.

2. Our perception of structure in these patterns raises questions as to the representation of that structure, the means by which that representation is computed from the image, and why the structure is perceived. Towards answering these questions, a psychophysical experiment was performed to determine the relationship between the perceived structure and the displacements between correlated dots. The experiment involved subjects judging patterns according to whether the dots appeared paired, and whether those pairs appeared to be locally parallel. The experimental variables were the type of pattern (e.g., radial, spiral), the dot density, and the displacement between corresponding dots.

3. For a given dot density, there is a critical displacement between corresponding dots such that the locally parallel structure is just perceptible. Over nearly a decade of dot densities the following rule applies: if more than two or three dots (2.3 on the average) lie closer to a given dot than its corresponding dot, then locally parallel structure among such dots cannot be perceived. This critical displacement is sensitive to the local dot density, making our perception of locally parallel structure in these patterns relatively independent of the density. This suggests that if a local computation is involved, that the size of the computational neighborhood is determined by the measured dot density in that locality.

4. It is argued that the perception arises from the local geometry. In particular, the parallelism between pairs of dots is taken to be fundamental. A representation of locally parallel structure is proposed: *virtual lines*. Each virtual line would represent the position, separation, and orientation between a pair of dots. A method for computing this representation is suggested: (1) virtual lines are constructed from every dot to each of the neighboring dots, and (2) those virtual lines that are locally parallel are selected.

5. An algorithm that embodies that method is described. The detection of parallelism is based on gathering local orientation statistics. An implementation of this algorithm is demonstated. Its performance parallels that observed experimentally.

6. It is demonstrated that the detection of local parallelism is not limited to dot patterns. The computation is apparently performed on *place-tokens*, distinguished points that have been abstracted from an image. The locally parallel structure implicit in textures such as fur, grass, or wood grain could be extracted by this method.

INTRODUCTION

A Moiré effect can be seen in patterns constructed by superimposing two copies of a random dot pattern where one copy had undergone some composition of expansion, translation, or rotation tranformations (figure 1a-1d) [Glass, 1969]. Our perception of structure in these "Glass patterns" has been taken as evidence that the visual system performs local autocorrelations [Glass, 1969; Glass and Switkes, 1976]. That is, the Moiré effect is due to the detection of pairs of correlated dots, each pair consisting of a dot in the initial pattern and the corresponding dot in the transformed copy.

Glass [1969] observed that the Moiré effect diminishes in the rotation-generated patterns as the amount of rotation increases. The periphery of the pattern (where the rotation causes the largest displacements) is the first to lose the circular organization. With sufficient rotation, one is left with an apparently random dot pattern. Furthermore, the Moiré effect will disappear if all but a small portion of the pattern is occluded [Glass and Perez, 1973]. Thus the effect is somehow dependent on the displacements between correlated dots, and the number of pairs of dots presented. The correlated dots need not be nearest neighbors for the effect to occur [Glass and Perez, 1973]. Recently it was shown that the pairs of correlated dots must correlate well in terms of orientation [Glass & Switkes, 1976].

This raises a number of interesting questions concerning (1) the representation of the perceived structure, (2) the means by which this representation is computed, and (3) why this structure is perceived. Prior to addressing these questions, the relationship between the perceived Moiré effect and the displacements between corresponding dots will be studied. From this, it will be concluded that the effect is derived from locally computed structure. Then, a method is introduced for computing a representation of this structure. Finally, a use for this representation is suggested.

EXPERIMENT

The experiment studied the effect of increasing the displacement between corresponding dots on the Moiré percept. The goal was to determine the maximum tolerable displacement as a function of the dot density.

Method

Glass patterns

The patterns consist of two superimposed copies of an initial dot pattern. Glass and Perez [1973] used random dot patterns. However, the use of random dot patterns confounds the Moiré effect with clusters, sparse regions, and especially, chains of dots.

Figure 1. Glass patterns constructed from a pseudo-random dot pattern and a superimposed copy of that pattern which has undergone some homogeneous displacement transformation. The patterns contain approximately 800 dots (ρ= .0124). The translation, spiral, radial, and concentric patterns (figures la-ld) all have displacements between corresponding dots of 7.7 units (pattern dimensions are 256 by 256 units), resulting in 1.95 neighbors lying nearer to a given dot than its corresponding dot. Figure le is a composite pattern composed of portions of the patterns in figures la-ld. The local structure is seen to be independent of the global organization (see discussion, p. 10). In figure lf, the corresponding dots are separated by 10.0 units, thus 3.75 extraneous neighbors lie closer to a given dot than the corresponding dot. While the corresponding pairs of dots are not seen as paired, nor as locally parallel, the pattern still appears radial. The radial effect is due to chains and clusters in the initial dot pattern that are selectively enhanced when the transformed copy is superimposed (see discussion, p. 5). Since these effects persist for very large displacements, the patterns used in the experiment were constructed in a manner that avoids the formation of chains.
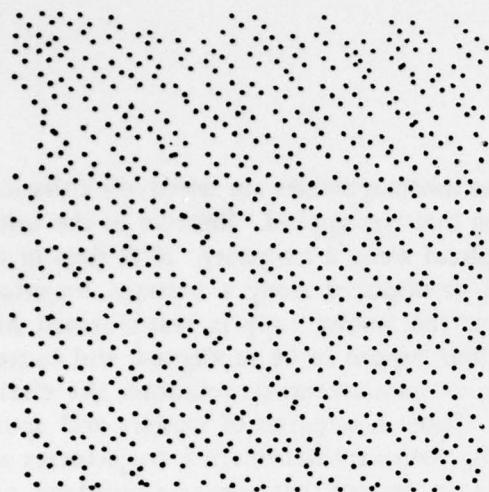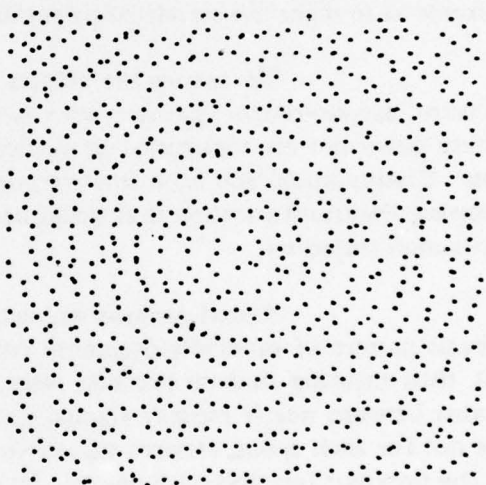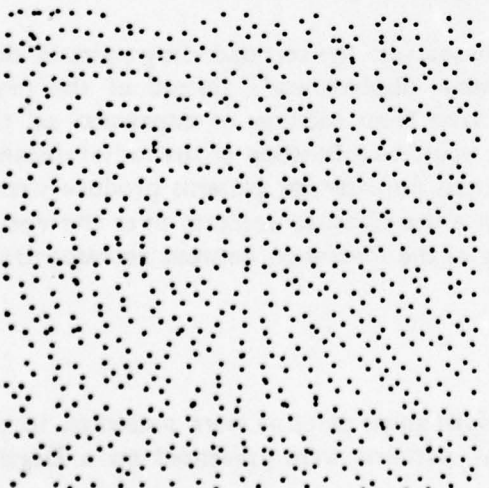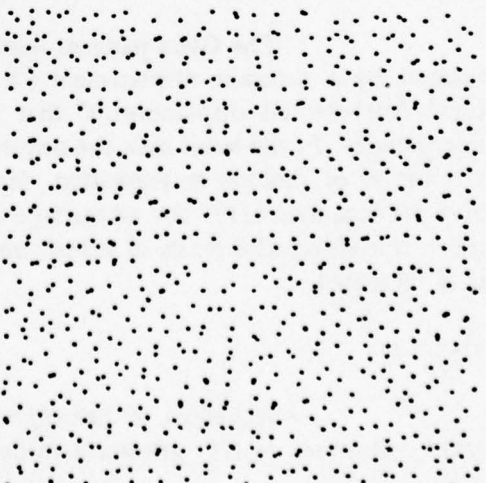
A
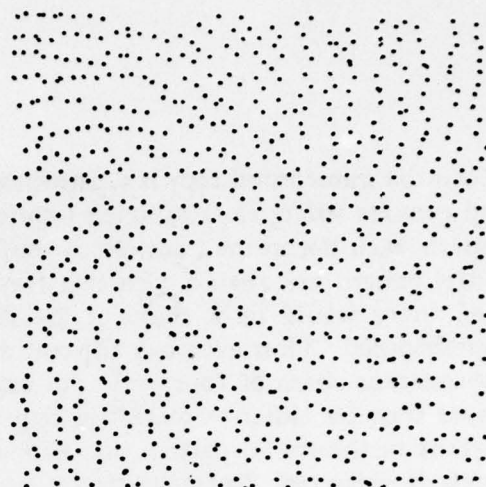
B
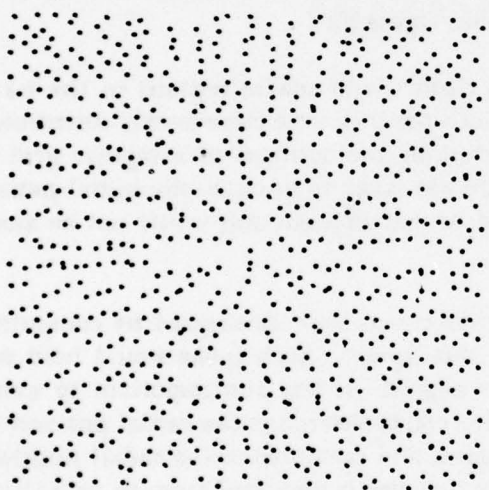
C

D

E

F

When the transformed copy is superimposed, these inhomogeneities are selectively enhanced, and provide strong clues as to the transformation that was applied. Relative to the initial pattern, each dot in the transformed copy is displaced along a trajectory. If N dots in the initial pattern are aligned such that they would be displaced along a common trajectory, then there would be a chain of 2N dots after the second copy is transformed and superimposed. Thus even two adjacent dots, if they happen to be so aligned, will cause a conspicuous chain of four dots. For expansion or rotation transformations, the chains would then be radial or concentric, respectively. Those boundaries of clusters and sparse regions in the initial pattern that happen to align with transformation trajectories are similarly enhanced. Consequently, clusters and sparse regions that appear amorphous and randomly oriented in the basis pattern appear wedge shaped in radial Glass patterns, or crescent shaped in rotational patterns. These clues persist when the transformation is so extreme as to make the correlated pairs indiscernible (figure 1f).

To reduce the effects due to clusters and sparse regions in the basis patterns, pseudo-random patterns were used in which the dots were more evenly distributed. These patterns were constructed by randomly perturbing the positions of a regular grid of dots. Chains would still arise, however, unless care was taken to generate the initial pattern knowing the transformation that would be applied, so that adjacent dots would not lie along a common trajectory.

Radial patterns without subjective chains were constructed by computing a basis pattern of randomly positioned dots on virtual spokes. Each spoke would hold one dot, thus insuring that no two dots were radially aligned. It was also important to avoid chains between nearly radially aligned dots. Therefore, to determine the radial position of the dot for each spoke, random values were computed and compared to the radial positions of the previous few dots until one was found to be sufficiently separated from its neighbors. The minimum allowed separation and the number of prior dots to be examined were empirically chosen so that the Glass pattern presented no subjective chains.

The Glass patterns were constructed with the corresponding pairs of dots separated by a constant displacement ("homogeneous displacement"), instead of the more natural "differential displacements" that would arise from rotation or expansion of the whole pattern. In the latter case, the displacement would be a function of the radial distance to the center of rotation or expansion. Homogeneous displacement patterns produce strong Moiré effects, and offer the advantage that since the effect is uniform over the entire pattern, the effect also tends to vanish uniformly as the separation between corresponding dots is increased.

Presentation

Sequences of homogeneous displacement patterns were presented to six unpaid volunteer M.I.T. graduate students. All patterns were presented on a Digital

Equipment Corporation GT-44 CRT display in a darkened room on a 23.5 by 23.5 cm. screen from a distance of 115 cms. (11.5 degree visual angle).

In the following, the dot density, ρ = (number of dots in pattern) / 256^2. The first series of presentations consisted of chainless radial patterns of five dot densities ranging from ρ=.00298 (195 dots) to ρ=.00884 (580 dots). For each dot density, 8-10 patterns were constructed with a range of displacements (between corresponding dots) for which the Moiré effect ranged from obvious to inapparent. A total of 45 patterns were presented in randomized order, in three sequences of 15 patterns each. Each sequence was viewed three times by each S, with the S instructed to judge each pattern numerically: "0" if the pattern appeared unstructured, "1" if the dots appeared to be paired, "2" if the pairings were locally parallel (i.e., while fixating a pair of dots, the neighboring dots also appeared paired and aligned with the fixated pair), and "3" if the parallelism appeared particularly strong. They were encouraged to sample several places on each pattern (avoiding the center and extreme periphery) before making their judgement, and to interpolate between these values according to the appearance in those localities. The presentation time was open ended, however Ss usually took 3-5 seconds per judgement.

A second series of presentations consisted of very low density patterns (ρ=.00096, 65 dots). Four types of patterns were used (radial, concentric, spiral, and translation). For each type, seven patterns of differing dot displacements provided obvious to inapparent Moiré effects. The 28 patterns were presented in randomized order as a single sequence. The sequence was presented three times to each S, and the S was asked to judge the patterns in the same manner as before, and to name the type of pattern as well. A typical response would have been "1.6 R" meaning "the dots appear paired, moreover in most places the pairings appear aligned; the overall pattern is radial."

Results

The responses of each S were separately tabulated, and for each sequence, that *critical displacement* for which the locally parallel pairings were just perceptible (i.e., an interpolated judgement of 1.5) was determined. The mean critical displacement for each density was then computed (see figure 2a). The data in figure 2a can also be expressed as follows: Define D to be the displacement between corresponding dots (constant across the pattern). Then consider a circular neighborhood of radius D centered on any given dot. The corresponding dot lies somewhere on the circumference of that circle. The number of other dots that would be expected in that neighborhood (i.e., to lie closer to the given dot than its corresponding dot) is a function of the dot density, specifically

$$N = \rho \pi D^2$$

Figure 2b shows a plot of N versus density computed from the averaged critical displacements of figure 2a. The mean N values for radial patterns were 2.31 (ρ=.00096), 1.91 (ρ=.00298), 2.33 (ρ=.00443), 2.37 (ρ=.00587), 2.36 (ρ=.00739), and 2.36 (ρ=.00884). The mean for ρ=.00298 is significantly less than the other means, as indicated by a t-test (p<0.05, t=2.83,

Figure 2. For a given dot density, there is a critical displacement (between corresponding dots) beyond which pairings between these dots cannot be perceived. This critical displacement (associated with an interpolated judgement of 1.5) was determined for each S, for each density. In figure 2a, the mean of these critical displacements is plotted as a function of dot density. For $\rho=.00096$, the critical displacement was determined for translational (T), radial (R), concentric (C), and spiral (S) Glass patterns that were constructed from a pseudo-random dot pattern (in order to minimize clustering effects). The higher density patterns were all *chainless radial* patterns.

If the critical displacement is taken as the radius of a circle, then for a given dot density, one can determine the number of dots that lie nearer to a given dot than its corresponding dot, when the pairings are just perceptible. Each vertical bar indicates two standard deviations. Whatever computation we perform on these patterns, it is relatively independent of the dot density. If local computations are performed over small neighborhoods (as will be argued), then this result suggests that the size of the neighborhood is determined by the measured dot density.
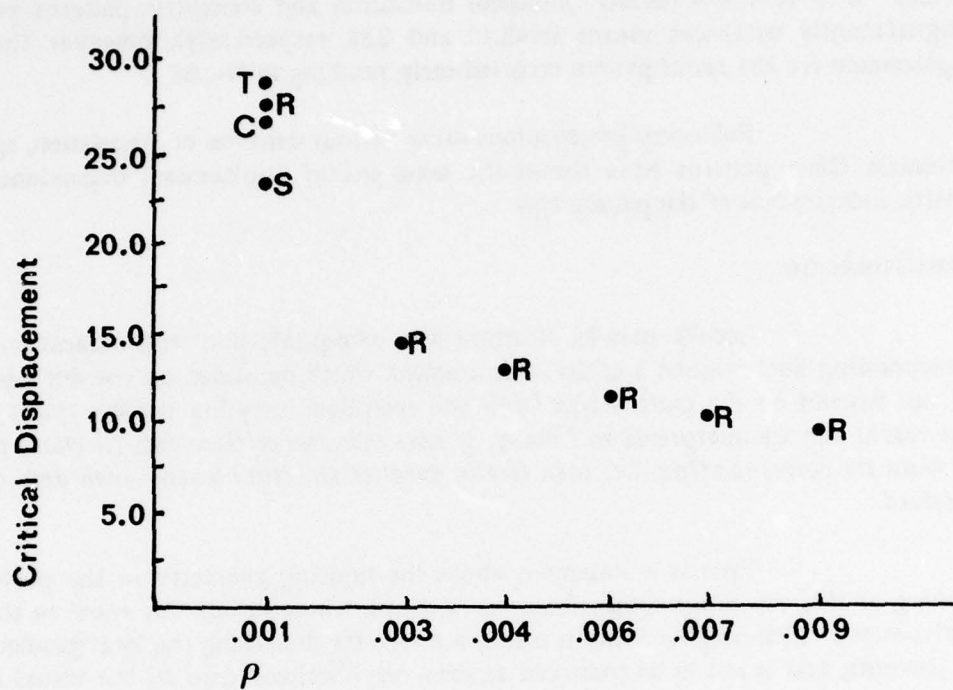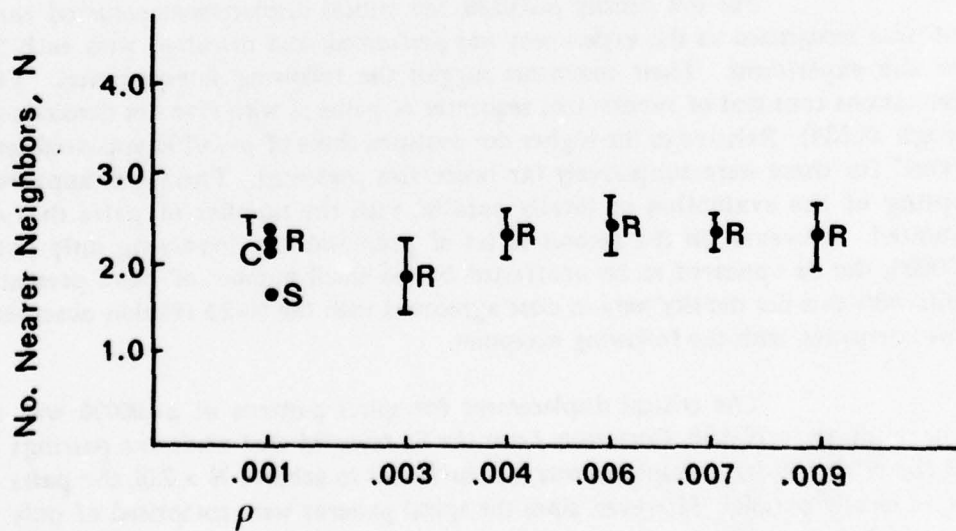
Figure 2a

Figure 2b

d.f.=32). The very low density (ρ=.00096) translation and concentric patterns resulted in insignificantly different means (N=2.40 and 2.31, respectively), however the critical displacement for the *spiral* pattern occurred early, resulting in N=1.68.

Follow-up presentations using various densities of translation, spiral, and concentric Glass patterns have shown the same critical displacement dependency on dot density, independent of the pattern type.

Conclusions

Locally parallel structure was perceptible until the separation between corresponding dots reached a critical displacement, which depended on the dot density, and did not depend on the pattern type (with one exception: very low density spiral patterns). The results can be interpreted as follows: *if more than two or three dots lie closer to a given dot than its corresponding dot, then locally parallel structure among such dots cannot be perceived.*

This is a statement about the limiting geometry in the patterns. In arriving at this result, a neighborhood was defined, whose radius was equal to the critical displacement. This neighborhood is merely a means for describing the local geometry of the dot patterns, and is not to be construed as some neighborhood used by the visual system in perceiving these patterns. Later, a computational neighborhood will be introduced.

For dot density ρ=.00298, the critical displacement occurred early. This trend was recognized as the experiment was performed, and discussed with each S directly after the experiment. Their comments suggest the following interpretation. The initial presentations consisted of randomized sequences of patterns with five dot densities (ρ=.00298 through .00884). Relative to the higher dot densities, those of ρ=.00298 appeared less "locally parallel" for there were subjectively far fewer dots presented. There was apparently some coupling of the evaluation of locally parallel with the number of pairs that could be evaluated. However, in the second series of presentations, involving only patterns of ρ=.00096, the Ss appeared to be unaffected by the small number of pairs presented. The results with this dot density were in close agreement with the N=2.3 relation observed for the higher densities, with the following exception.

The critical displacement for spiral patterns of ρ=.00096 was relatively small, resulting in N=1.68. Comments from the Ss revealed that while the pairings could be held for relatively large displacements (i.e., sufficient to achieve N > 2.0), the pairs were not seen as locally parallel. However, since the spiral patterns were comprised of only thirty or so pairs of dots scattered over the display, one would not expect the widely separated pairs to appear locally parallel. At least with concentric and radial patterns some of the neighboring pairs relative to a given pair will be parallel. For example, with a concentric pattern, those pairs that lie on the same (or a nearby) radius will be approximately parallel.

In fact, the results with very low density radial, translation and concentric patterns were similar, and in close agreement with the results from higher density patterns.

The critical displacement is sensitive to the *local* dot density, for if a Glass pattern is constructed with varying dot density but constant displacement between corresponding dots, the effect is apparent only in those neighborhoods where N would be less than two or three.

Whatever computation we perform on these patterns, it is relatively independent of the actual dot density. If a local computation is involved, then the angular extent of the neighborhood is determined by the measured dot density in that locality. Before arguing that the computation is local, one further result should be mentioned.

There had not been any investigation into the time required to perceive the Moiré effect in these dot patterns. In fact, it was not known whether eye movements are necessary for developing the impression of structure. To study this, a sequence of masking random dot patterns were presented before and after a single Glass pattern. The eight masking patterns had the same dot density as the Glass pattern, and the sequence was presented without pauses between frames. The frame rate was the experimental variable. It was assumed that in order to detect the Moiré effect, that the locally parallel structure would have to be determined within the time that the Glass pattern was presented. Thus the minimum presentation time would approximate the minimum computation time for determining the locally parallel structure. Note that once the local structure is determined, the global Moiré effect may continue to develop during the presentation of the subsequent masking patterns. It was found that at 80-90 msec/frame, one could reliably name the type of pattern. At 100-110 msec/frame one could name two different Glass patterns that were presented in succession while embedded in the masking sequence. Since the two patterns were presented in the same visual region, it is more likely that we perform two fast computations in sequence rather than two slower ones in parallel. Thus the computation of locally parallel structure is relatively fast, and does not require eye movements.

Global impressions derived from locally parallel structure

Glass [1969] suggested that in our perception of these patterns, local correlations from different regions of the visual field are combined to form a simple global percept. That is, the processing is bottom-up, in contrast to the top-down alternative in which the overall structure is somehow determined, and that in turn influences the local percept. To support the bottom-up hypothesis, a composite Glass pattern (figure 1e) was created from portions of figures 1a through 1d. If the overall organization were to influence the perceived local structure, then one would expect that a neighborhood of dots taken from one pattern and embedded in another would appear differently in its new surroundings. However, the Moiré effect in any locality of figure 1e appears as it does in the original pattern (except along the boundaries where the neighborhoods have changed). The new

global geometry does not influence the local structure.

It is easy to demonstrate that the Moiré effect requires a number of dots in order to be seen. If one masks out progressively greater portions of the pattern, the effect diminishes until so few dots are left that one becomes aware of coincidental arrangements among those dots [Glass & Perez, 1973]. If, however, the pattern is initially masked except for a few dots, and progressively larger neighborhoods centered on the initially visible dots are revealed, then the initial, coincidental groupings of dots are replaced by pairwise groupings. As the pattern becomes more fully exposed, those pairings remain, and are seen to be locally parallel. When our awareness is on the overall pattern, we see a Moiré effect, while under scrutiny, we see pairs of dots. Note that very close pairs of dots can also be seen that are oriented contrary to the Moiré structure in that vicinity.

Thus two subjective impressions can be studied: the Moiré effect, and the pairings of dots. Consider a pattern consisting of a large number of dots, where the corresponding dots are nearest neighbors (figure 3a). Both impressions are strong, however the prevailing impression is one of global structure. If the displacements are increased (holding the density constant), then the pairings become less obvious, while the Moiré effect is still strong (figure 3b). However, the effect is weak when the density is small, regardless of whether the corresponding dots are nearest neighbors and the pairings are strong (figure 3c) or not (figure 3d).

It is hypothesized that the global structure (e.g., "spiral", "radial") is derived from the local pairings, and constitutes a later, distinct computational problem. Thus this paper is directed towards the more fundamental problem, how the pairings are represented, and how that representation is computed.

Representing locally parallel structure

A natural representation for a perceived local pairing would be a *virtual line*. Each virtual line would represent the position, separation, and orientation between a pair of dots. The proposed representation of the local structure is simple, being a discrete, spatial arrangement of virtual lines. The Moiré effect would then arise from this local structure. The strength of the effect would be dependent on the size of the population, the length of the virtual lines, and their collective geometry.

The orientation of the local structure is represented only at discrete points in the image. Would a continuous representation be necessary? Consider an analogy to the representation of depth from stereopsis. Discrete stereo disparity clues result in a perceived surface that is continuous (e.g., in random dot stereograms [Julesz, 1971]). The strong impression of depth that we assign to all points in the image suggests that underlying this percept is a continuous representation of depth. However, a continuous representation for locally parallel structure would not be appropriate, for there is no evidence that we attribute

a sense of orientation to all points in the pattern.

Computing the representation

The fundamental problem in computing the representation is to determine which groupings to construct, for in the vicinity of any dot there are many neighboring dots with which the given dot can be paired. We understand that the perceived pairings are between corresponding dots, and that these pairings are seen to be locally parallel. While the corresponding dots cannot be known *a priori*, the virtual lines that would connect them would be locally parallel. Therefore it is hypothesized that the following method underlies the computation:

(1) virtual lines are constructed from every dot to each of the neighboring dots, and
(2) those virtual lines that are locally parallel are selected.

Constructing virtual lines

The first step is to construct the virtual lines that radiate from each dot to each of its neighboring dots. This raises a question as to how large the neighborhood centered on each dot should be. Since the computational problem is to select one virtual line (that which extends to the corresponding dot) from each neighborhood, it would be optimal to have the neighborhood just large enough to include its corresponding dot. A larger neighborhood would merely include more extraneous dots, a smaller one would fail to take the corresponding dot into consideration. Since there is no *a priori* knowledge of the position of the corresponding dot for any given dot, that neighborhood should be roughly circular.

The demonstrated independence of the Moiré effect from the angular extent of the pattern suggests that the neighborhood radius is a function of the local dot density. For now, consider that a neighborhood is defined on the basis of the local dot density, and that it is large enough to hold a few nearby dots. Better insight into the size of the neighborhood will be provided by the performance of an implementation.

Representing a small number of virtual lines that radiate from the center of the neighborhood poses no significant computational problems. In the proposed algorithm, a virtual line is represented by two quantities, an orientation, and a weighting. The weighting is greater for shorter lines, resulting in an algorithm that favors nearer pairings. This will be discussed in more detail later.

Selecting the locally parallel lines

Given the virtual lines, the problem is now to extract those that are locally parallel. This problem can be solved simultaneously for each dot: that virtual line

Figure 3. There are two subjective impressions that can be studied in these patterns, the Moiré effect, and the pairing of corresponding dots which gives rise to locally parallel structure. Figure 3a and 3b have high dot density ($\rho=.02$, 1300 dots) while figures 3c and 3d have low density ($\rho=.001$, 70 dots). With a high dot density, the Moiré effect appears strong regardless of whether the corresponding dots are nearest neighbors (figure 3a, N=.778) or not (figure 3b, N=2.25), while with a low dot density, the effect is weak in both cases (figure 3c, N=.377; figure 3d, N=2.06). In this study, the emphasis has been on the computation of locally parallel structure, under the assumption that the Moiré effect is derived *from* the local structure.
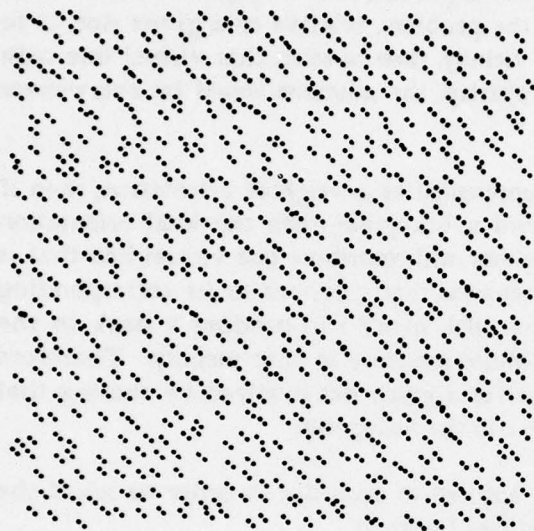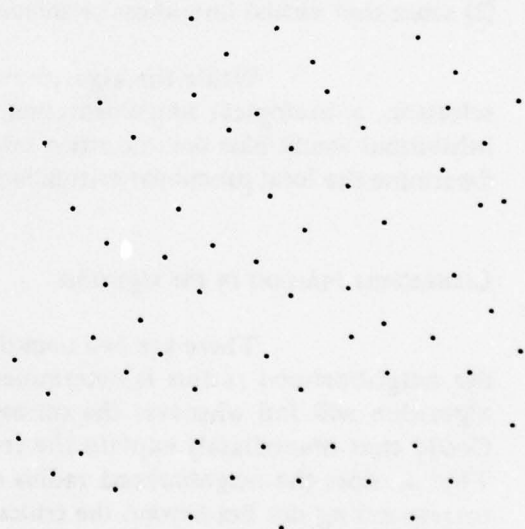
A
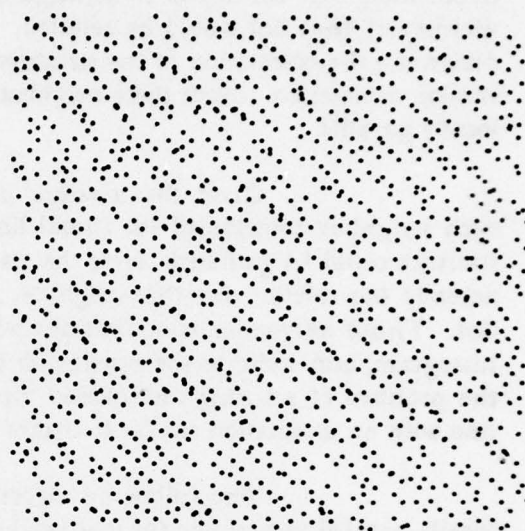
B
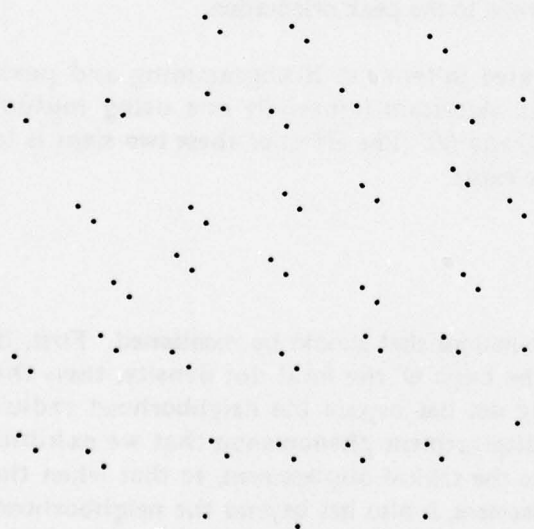
C

D

(from the given dot to one of its neighbors) which is parallel to the Moiré structure in the vicinity of that dot would be selected. Thus the problem, relative to a given dot, is to determine the orientation of the structure in its vicinity, then to select that virtual line with similar orientation. Since these neighborhoods overlap, the solutions would be everywhere locally parallel.

Given that a virtual line is represented as a weighted orientation, then if each neighbor contributed its virtual lines toward a histogram, then the local orientation statistics could be gathered. Note that each neighbor will contribute one virtual line that is actually the solution for that neighbor, i.e., it connects that neighbor to *its* corresponding dot. Those particular contributions will be parallel, hence will produce a peak in the histogram, and indicate the orientation of the Moiré structure in that vicinity. Therefore the problem of selecting the solution virtual line for a given dot is solved by chosing that line with an orientation similar to that of the peak in the histogram.

The following algorithm is applied to each dot in order to select the locally parallel virtual line for that neighborhood (see figure 4).

(1) histogram the orientations of the virtual lines of its neighbors,
(2) determine the peak orientation from the histogram, and
(3) select that virtual line whose orientation is closest to the peak orientation.

While the algorithm is phrased in terms of histogramming and peak selection, a biological implementation of this algorithm (especially one using mutual inhibition) would blur the distinction between (1) and (2). The effect of these two steps is to determine the local prominent orientation, if one exists.

Limitations inherent in the algorithm

There are two immediate limitations that should be mentioned. First, if the neighborhood radius is determined on the basis of the local dot density, then the algorithm will fail whenever the corresponding dot lies beyond the neighborhood radius. Could that immediately explain the critical displacement phenomenon that we exhibit? That is, does the neighborhood radius equal to the critical displacement, so that when the corresponding dot lies beyond the critical displacment, it also lies beyond the neighborhood radius, hence not considered by the algorithm? Probably not, for within the radius of the critical displacement there are only two or three neighbors, which would be an insufficient sampling from which to produce a histogram with a reliable peak.

The algorithm is also limited by the orientation resolution, both in the representation of the virtual lines, and in their summation into the histogram. To illustrate, suppose that each dot has N neighbors. Then the area under the peak in the histogram

would be at most N, while the total histogram area would be N^2, distributed over M "buckets" (determined by the orientation resolution). For any given M, if N is sufficiently large, the peak will be submerged in the histogram.

The Gestaltists recognized that we tend to see rectangular grids as either columns or rows, depending whether the vertical or horizontal spacings are smaller, respectively. The algorithm shares this behavior when proximity weighting is introduced. Without this proximity metric, the interior dots would have four strong peaks, corresponding to pairings in the principal diagonal orientations, the vertical, and the horizontal. Since the nearer pairing orientations are emphasized in the histogram, then that peak contributed by the nearer pairings is emphasized, allowing that orientation to be selected. However, proximity weighting will also limit the algorithm. Suppose that the displacement between corresponding dots is such that there are several extraneous nearer neighbors. The virtual lines to these dots would be emphasized more than the virtual line to the corresponding dot. As this would occur to the virtual lines in any vicinity, the contributions from the locally parallel lines would be relatively less effective in producing a peak in the histogram. Therefore, as the number of nearer neighbors increases (i.e., the displacement increases for a given dot density), the peak will become less significant.

If the neighborhood radius is large relative to the curvature of the structure (e.g., near the center of a radial or concentric pattern, especially with low dot densities), then the notion of "locally parallel" breaks down. The peak in the histogram would broaden, and selection of the solution orientation would become less reliable. The experiment demonstrated that locally parallel structure is difficult to perceive in low density dot patterns were the curvature is considerable.

In summary, the algorithm is fundamentally limited by three factors: the orientation resolution, the neighborhood size, and proximity weighting.

An implementation of the algorithm

An implementation in LISP has demonstrated that the algorithm is capable of computing the Moiré structure representation. The performance of the algorithm on various Glass patterns is demonstrated in figure 5, where the local orientation, as determined by the algorithm, is indicated by short line segments centered on the dots.

The virtual lines that radiate from a given dot to its neighbors were encoded by their orientations (the orientation resolution was 10 degrees) weighted in a simple manner by their length relative to the neighborhood radius, depending on whether the neighboring dot was nearer than a quarter, less than one half, or greater than half of the neighborhood radius. The weights were 1, 2/3, and 1/3, respectively.

Figure 4. The algorithm for computing locally parallel structure has three fundamental steps. Place tokens that are defined in the image are the input to the algorithm. The algorithm is applied in parallel to each place token. Since, in the case of the Glass patterns, each dot contibutes a place token, the first step is to construct a virtual line from that dot to each neighboring dot (within some neighborhood centered on the dot). A virtual line would represent the position, separation, and orientation between a pair of neighboring dots. To favor relatively nearer neighbors, relatively short virtual lines are emphasized. The second step is to histogram the virtual lines that were constructed for each of the neighbors. For example, the neighbor D would contribute virtual lines DA, DF, DG, and DH to the histogram. The final step (after smoothing the histogram) is to determine the orientation at which the histogram peaks, and to select that virtual line (AB) closest to that orientation as the solution.
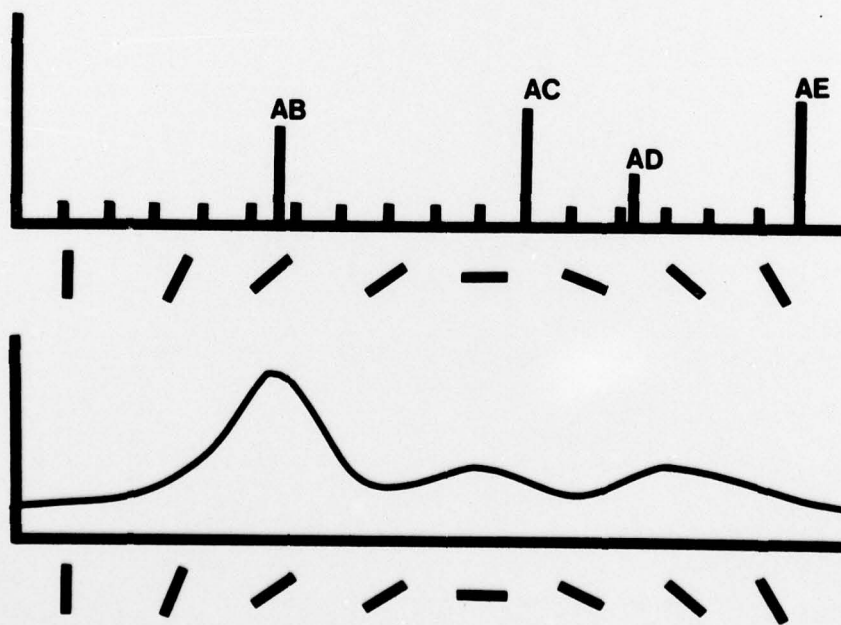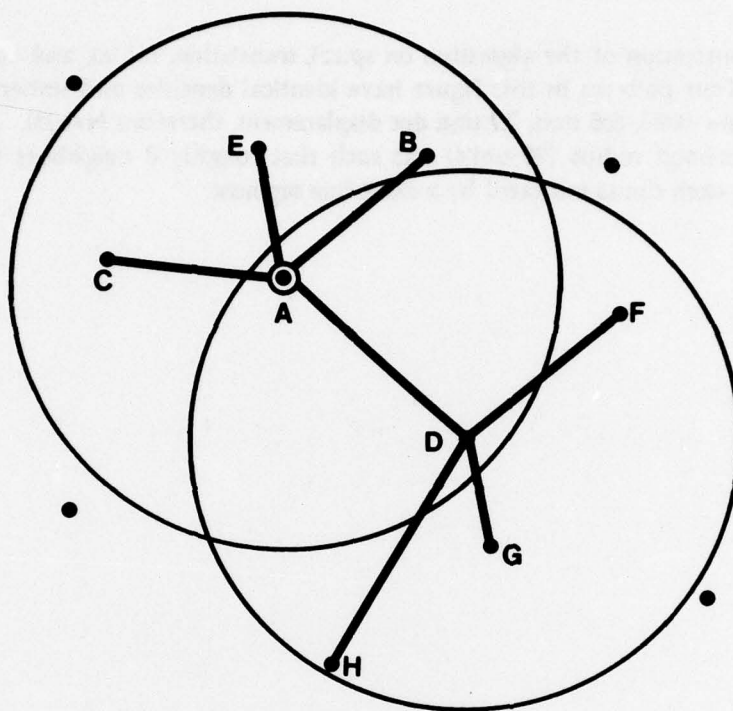
Figure 5. Demonstration of the algorithm on spiral, translation, radial, and concentric Glass patterns. The four patterns in this figure have identical densities and underwent the same displacements ($\rho = .0085$, 556 dots, 7.7 unit dot displacement, therefore N=1.33). The algorithm used a neighborhood radius (20 units) was such that roughly 8 neighbors were included. The solution at each dot is indicated by a short line segment.
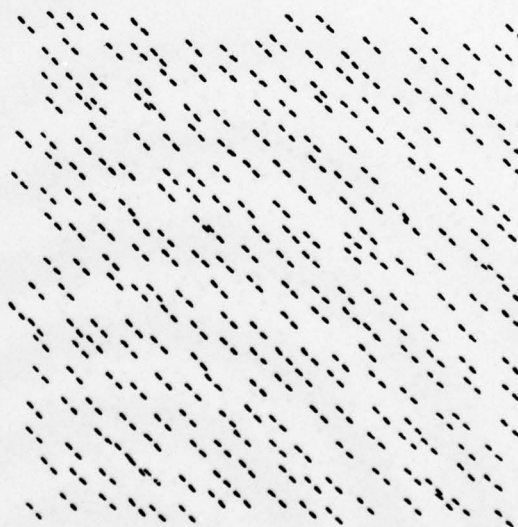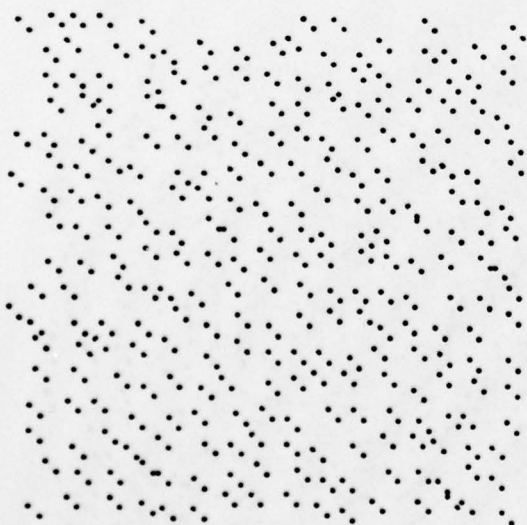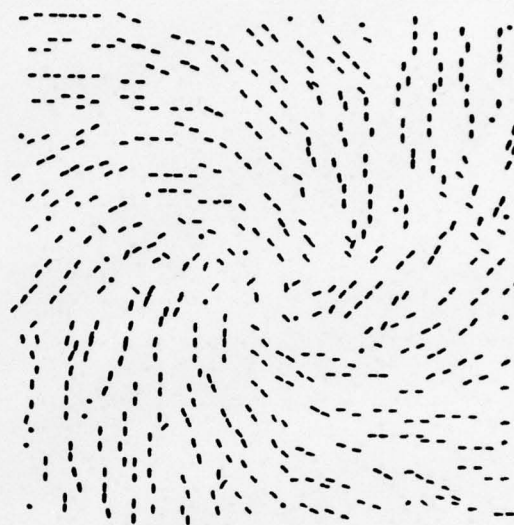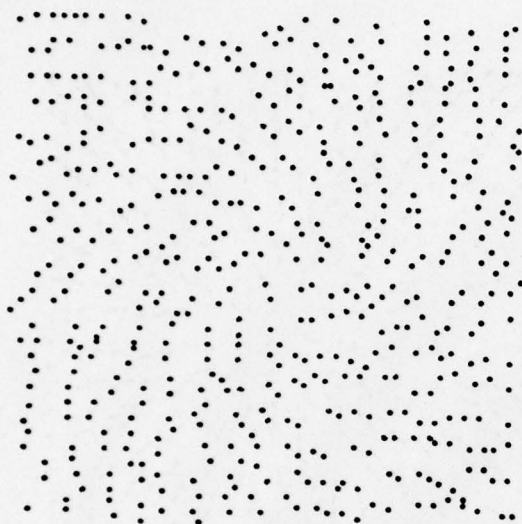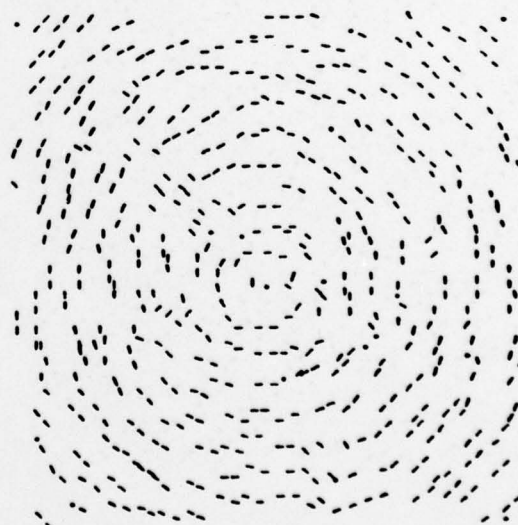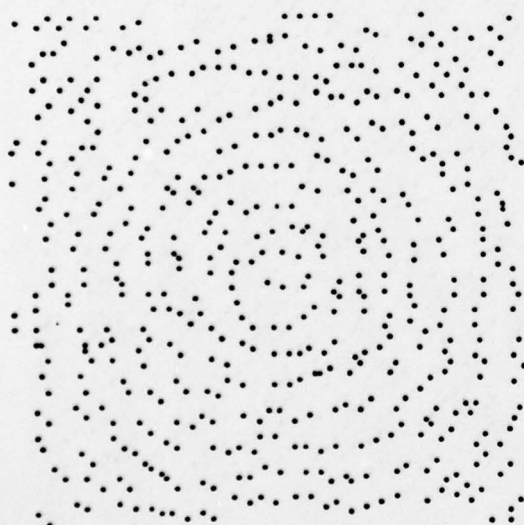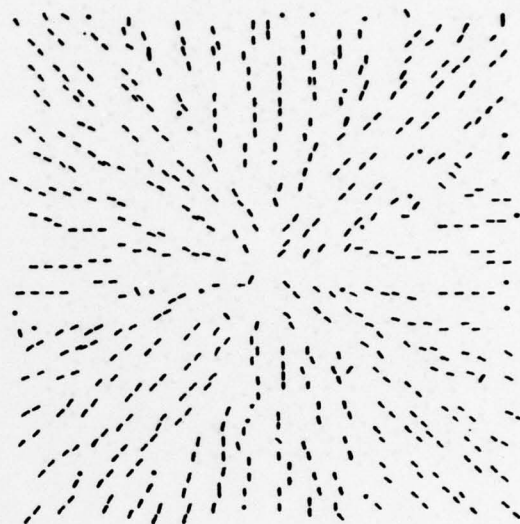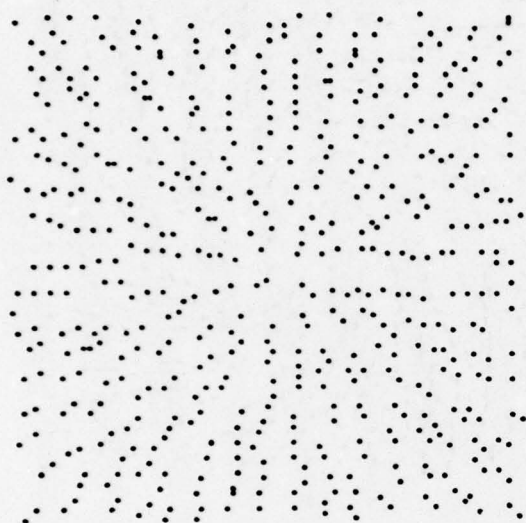
The second step was to determine the solution orientation relative to each dot, computed by histogramming the weighted orientations associated with each of its neighboring dots and determining the peak orientation. Various criteria were studied for determining the peak of the histogram, with the conclusion that since the total area under the histogram curve is small, stringent criteria that require that the the peak be "significant" would often not be satisfied. With the exception of translation Glass patterns, the structure would not be strictly parallel in any neighborhood, causing the few contributions to the peak to be scattered over several adjacent histogram "buckets". Therefore a smoothing operator was applied to the curve to accentuate the peak, and that orientation with the maximum value was selected.

The final step was the selection of the solution virtual line from the set associated with each dot. That line whose orientation was nearest to the peak orientation was chosen and displayed graphically. If no virtual line was within 15 degrees of the peak orientation, then a dot was displayed, signifying that no solution was found.

Insight into our critical displacment limitation?

If one were to accept the conjecture that we share the same algorithm for the perception of locally parallel structure, then could the LISP implementation provide us with insight into the cause for the observed limitations in our perception of the Moiré effect?

By varying the orientation resolution and neighborhood radius, the implementation of this algorithm can perform with either greater than or less than human ability (measured by the critical displacement between corresponding dots). An empirical study of this implementation was undertaken in order to determine if a particular choice of parameters would result in performance that closely matches ours. If that were found, then it would be interesting to reflect on the cause for the implementation's limitation given those parameters. Four orientation resolutions were used: 45, 33.3, 22.5, and 10 degrees (4, 6, 8, and 18 buckets). For each resolution, the algorithm was then run on translation and radial Glass patterns, while varying the neighborhood size. The first step was to increase the neighborhood size (measured by the number of included neighbors) until the performance was just breaking down at the critical displacement (N=2.36) while closely matching ours for lesser displacements. Then the algorithm was run (with the same neighborhood size) on radial patterns of various dot displacements, in order to verify that curvature does not effect the performance. It was found that reasonable performance could be achieved with as little as 33.3 degree orientation resolution when the neighborhood radius is such that only six or seven neighbors were included. This neighborhood radius is sufficiently small that curvature within that vicinity is insignificant, thus the performance is similar for radial and translation patterns.

The conclusion drawn from this is that the parameter that governs the

limiting performance is the neighborhood radius. Presumably, in choosing between (1) having a large sampling from which to make statistical decisions, and (2) restricting the area over which the samplings are taken, in order to avoid curvature, that the latter consideration is favored. The inevitable consequence then, is that the peak will often not be correctly distinguished from the noise. As discussed, proximity weighting helps when the corresponding dot is relatively nearby within the neighborhood, and hurts when it is near the perimeter of the neighborhood. When the corresponding dot is displaced by approximately 60 percent of the neighborhood radius (ratio of critical displacement to neighborhood radius) then the performance becomes significantly deteriorated.

While the performance is satisfactory with low orientation resolution, the performance with 10 degree resolution most closely parallels human performance. That is, if the solution line segments computed by the implementation do not correspond to the ideal solution in some small locality, it is often the case that we also perceive some anomolous groupings in that locality that are contrary to the overall Moiré structure. In summary, the algorithm exhibits human performance when the neighborhood is determined to be large enough to hold 6 or 7 neighbors, and the orientation resolution is 10 degrees.

How abstract are the virtual lines?

The proposed algorithm is based on virtual lines constructed between neighboring dots. The virtual line is an abstract construct that expresses a grouping between two elements in the image. Can a simpler explanation be found that would account for the Moiré effect, without having to construct some representation of groupings?

Glass [1969] suggested that the effect is evidence for local autocorrelation of the excitation of orientation-sensitive cortical units (presumably "simple cells" [Hubel and Wiesel, 1962]). According to this hypothesis, pairs of dots would tend to trigger these units when they happen to be aligned in their receptive fields. While the various coincidental pairings would result in the excitation of a large number of units, if their outputs were correlated over some neighborhood, the prominent orientation would correspond to the subjective flow orientation in that vicinity. Evidence that supports this hypothesis has been reported [Glass and Switkes, 1976].

However, there is some evidence to suggest that more is involved in our perception of parallelism in these patterns than simply the correlation of simple cell activity. Rival patterns will be described for which we prefer pairings between dots of similar intensity. Two consequences of this will be discussed: (1) that the Glass proposal does not correctly predict this preference, and that (2) we should consider the pairings as groupings between abstract places in an image.

Consider a Glass pattern constructed from the superposition of three patterns: an initial pattern, and two differently transformed copies. The resulting pattern is

potentially rivalrous, for there are two locally parallel structures (figure 6a). First consider the case when the dots are of equal intensity and the displacements undertaken by both transformations are equal. Locally parallel organization is difficult to perceive. However, with some effort we can extract either of the organizations, wherein the other (unpaired) dots are see as background.

Now, if the dots of the initial pattern and those of one of the transformed copies are displayed with low intensity, while the dots of the other transformed copy are of higher intensity, then we favor the organization consisting of pairings between low intensity dots. The subjective impression is one of a faint Moiré effect and a superimposed random pattern of bright dots. It is difficult if not impossible to perceive pairings between faint and bright dots as being locally parallel. If one fixates on such a pair, then the vicinity appears heterogeneous (i.e., to consist of *pairs* of faint dots mixed with individual bright dots).

The display apparatus gives us the facility to continuously vary the relative intensities of these two populations of dots. The display instructions specify two intensity levels, however, a potentiometer that governs the overall brightness can, in one extreme, make both intensity levels appear equally bright, while towards the other extreme make the lower intensity level effectively invisible while the higher level is still faintly visible. Thus all intensity ratios from 0:1 to 1:1 can be achieved. The rivalrous patterns appear ambiguous in the equal-intensity extreme (as in figure 6a). If one reduces the overall brightness, the lower-intensity dots become distinguishable from the higher-intensity dots, and pairings between the former are favored. In the extreme, these dots are so faint as to be insignificant, the brighter dots dominate, and the pattern appears random. At no point is there a preference for pairings between dots of differing intensity over those of like intensity.

It is difficult to account for this behavior with the mechanism based on correlated simple cell excitation. On the contrary, that proposal would predict the correlation to be stronger between faint-bright pairings, for units aligned with those pairings would be more excited that those oriented with the faint-faint pairings. What of the possibility that the faint-bright pairings do not enter into the correlation? One has merely to remove the competing faint dots in order to perceive a strong Moiré effect between the faint dots of the initial pattern and the bright dots of the remaining transformed copy.

It appears that some notion of similarity must be introduced into both proposed mechanisms. With the Glass proposal, the correlation must be on brightness as well as orientation and displacement (this may be difficult to provide with simple cells). Similarly, the histogram-based computation must introduce some notion of similarity. Clearly, one could introduce it in the same manner as proximity weighting (i.e., just as proximate dots are favored, so are dots of similar intensity). Then the virtual lines would

Figure 6. A rivalrous pattern (figure 6a) is created by superimposing two differently transformed copies. While there are two locally parallel structures in this figure, they are difficult to perceive. However, if the pattern is displayed on a CRT, with the dots of one of the two superimposed copies brighter than those of the other two copies, then the fainter dots produce a Moiré effect. This tendency to pair similar dots would not be expected if the local orientation is derived from correlated simple cell activity. A simple cell whose receptive fields holds a bright dot and a faint dot would respond more vigorously than one that holds a pair of faint dots.

                 In figures 6b-6d, a spiral Moiré effect is evident although derived from pairings between dots and short line segments. The lines are randomly oriented in figure 6b, while in figures 6c and 6d, the lines have global radial and translation organization, respectively.
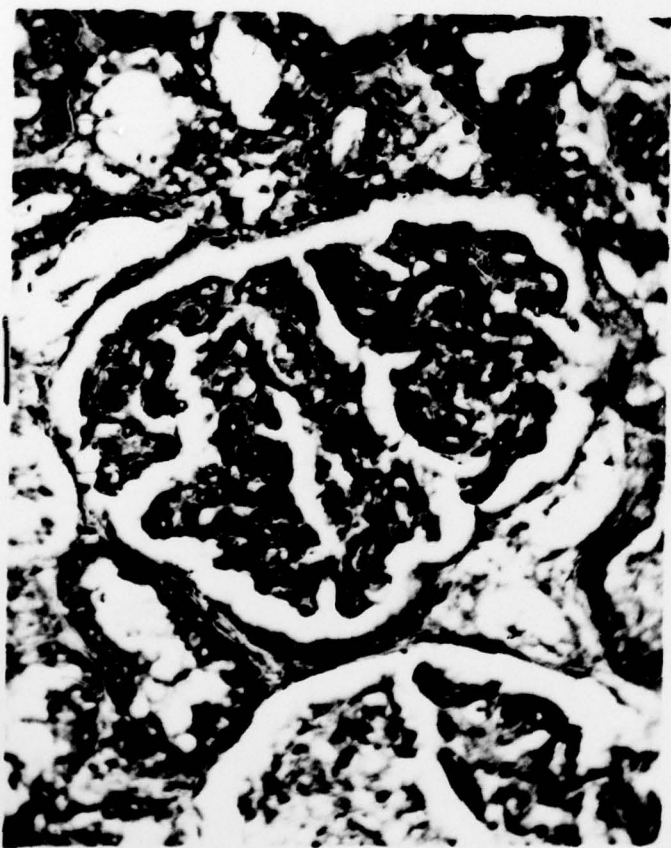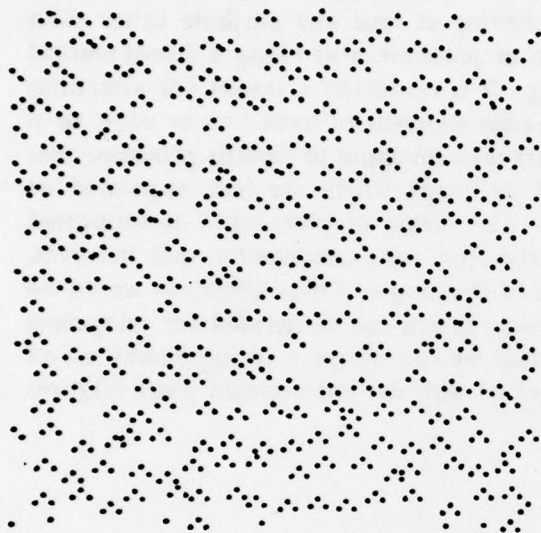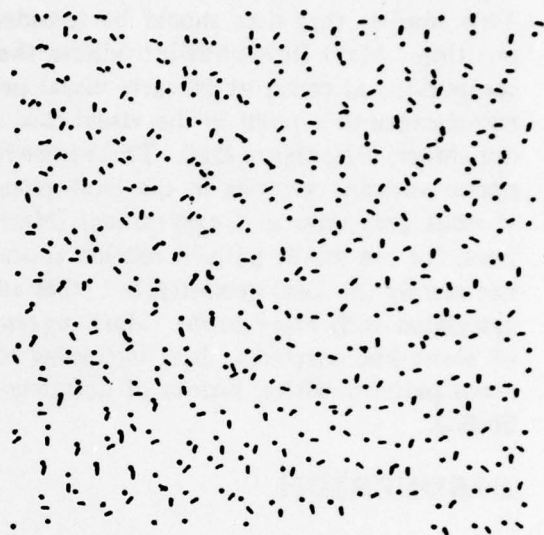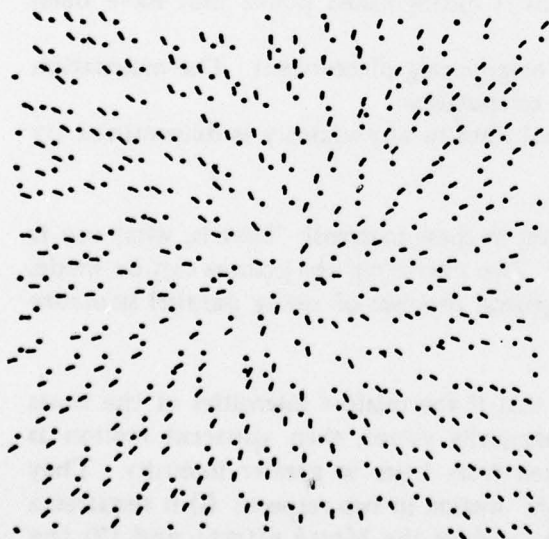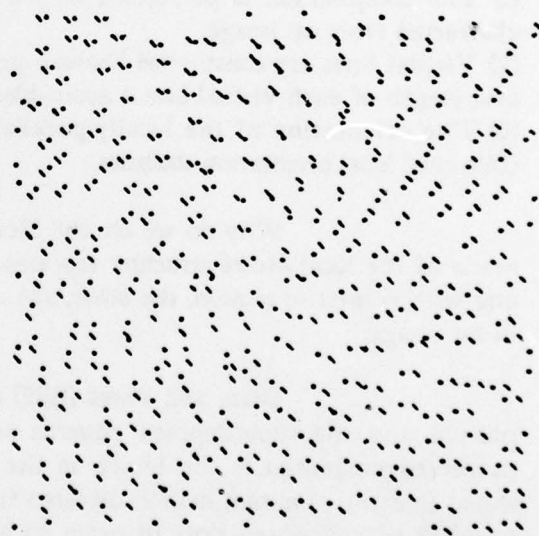
express three quantities: the orientation, separation, and similarity between a pair of dots. This implies that dots should be considered as having at least one attribute other than position. Marr [1976] has introduced the notion of *place-token* as being a fundamental computational construct in early visual processing. It is essentially a means for attaching significance to a point in the visual field (such as the endpoint of some line or edge, or a dot [Marr, 1976; figure 12a]). These place-tokens are then the input to various processes that notice various relations in the local geometry of an image, which are then expressed as various *groupings* and *aggregations* [Marr, 1976]. The notion of place-token is supported here, for the locally parallel relation appears to arise from some computation that involves, not merely the local geometry, but other attributes of the image. These attributes would be associated with place-tokens. Marr suggests that place-tokens can be defined for midpoints of short line segments. It is interesting to note that we can derive a strong Moiré effect from patterns where, instead of dots, one is presented with dot-line segment pairs (figures 6b-6d).

DISCUSSION

A representation of locally parallel structure has been shown to be amenable to a particularly simple computation. The following issues have been illustrated:

(1) The computation is performed on place-tokens -- distinguished points that have been abstracted from an image.
(2) Virtual lines are constructed between pairs of neighboring place-tokens. The orientation and length of each virtual line is accessible to the computation.
(3) The orientation of the locally parallel virtual lines in any vicinity is determined by collecting local orientation statistics.

Why do we see the Moiré effect in these patterns? That is, what use is made of the local Moiré structure representation? Two interesting conjectures can be made, one with respect to motion, the other, about the general problem of seeing parallel structure in an image.

Glass and Perez [1969] found that if the relative intensities of the basis pattern and the superimposed patterns are dynamically varied, then apparent motion is perceived tangential to the Moiré, in the direction from lesser to greater intensity. They noted that the apparent motion differed from "phi" motion in two respects: (1) it requires a number of correlated dots in order to be seen (as does the Moiré effect), and (2) the corresponding pairs of dots must be simultaneously (rather than alternately) presented. If the Moiré representation were involved with motion, it would be useful for expressing correspondence relations between successive images. For example, if the initial pattern and the normally superimposed pattern are shown in succession, apparent motion can be seen. For this to occur, we must be establishing a 1-1 correspondence between dots seen in the first and second images. The proposed virtual line representation would then express this
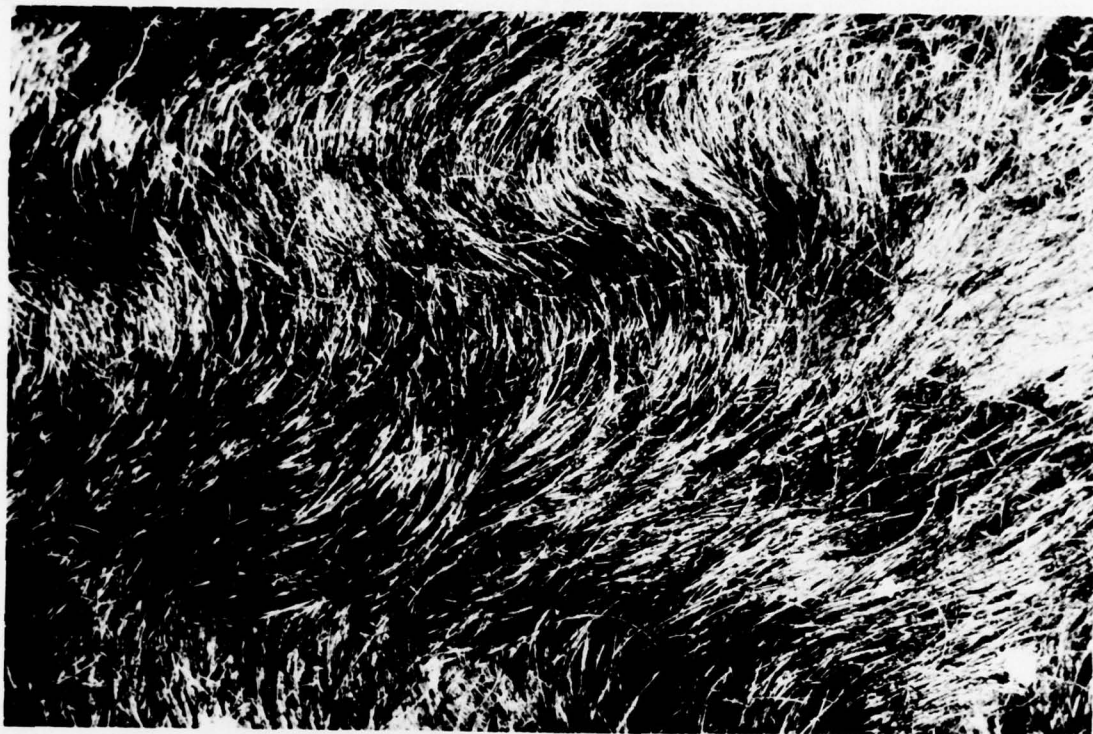
Figure 7. This photograph of human hair appears homogeneous, yet if analyzed in terms of intensity, line length, or line orientation, the image would be heterogeneous. It is suggested that the homogeneity that we perceive is the locally parallel structure. This structure could be extracted by a method based on computing local orientation statistics, and selecting those lines and edges that are parallel to the prominent orientation in the vicinity.

correspondence. The correspondence would be computed wholly on detected locally parallel trajectories.

The hypothesis that this algorithm computes the locally parallel structure that expresses motion correspondence is weakened by the observation that the algorithm, while sufficient for the Glass patterns, is insufficient for pairing corresponding dots between frames of dot patterns, when the displacements undertaken by the individual dots between frames is considerable. As discussed, the algorithm tends to fail if more than roughly three extraneous neighboring dots lie closer to a given dot than its corresponding dot. However, if the two patterns that comprise a Glass pattern are presented in succession, then we can perceive rigid motion when an order of magnitude more extraneous dots (greater than 40) lie closer than the corresponding dot. To account for this ability, an algorithm based on histogramming would require very fine orientation resolution in order to detect the peak. It is probably unreasonable to expect that fine of orientation resolution in early vision. Furthermore, the emphasis placed on proximate neighbors, which is evident in the Moiré effect, is not apparent in the apparent motion effect (the dots appear to move as if attached to a rigid invisible surface, in spite of very near neighbors). A computation based wholly on the local geometry, as is this algorithm, would probably not be sufficiently constrained to solve this motion correspondence problem. Temporal and other constraints must be incorporated as well [Ullman, 1977].

The second conjecture concerns the perception of locally parallel structure in an image. According to this hypothesis, Glass patterns present stimuli to processes that (1) define place-tokens in the image, (2) construct virtual lines between neighboring tokens, and (3) extract those that are locally parallel. The algorithm by which (3) is accomplished is presumably applicable to "actual" lines and edges as well. Natural images often contain locally parallel textures (e.g., fur, grass, wood grain), which would result in large numbers of parallel line and edge elements in a description of that image. This structure could be extracted by a method based on computing local orientation statistics, and selecting those lines and edges that are parallel to the prominent orientation in the vicinity. In figure 7 we perceive a certain homogeneity -- not of brightness, orientation, or line length -- but rather, of structure. That structure is locally parallel.

References

Glass, L. 1969 Moiré effect from random dots. *Nature 243*, 578-580.

Glass, L. & Perez, R. 1973 Perception of random dot interference patterns. *Nature 246*, 360-362.

Glass, L. & Switkes, E. 1976 Pattern recognition in humans: correlations which cannot be perceived. *Perception 5*, 67-72.

Hubel, D.H. & Wiesel, T.N. 1962 Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.) 160*, 106-154.

Marr, D. 1976 Early processing of visual informantion. *Phil. Trans. Roy. Soc. Lond. B. 275*, 483-519. (also available as *M.I.T. A.I. Lab. Memo 340*).

Ullman, Shimon 1977 The interpretation of visual motion. *M.I.T. A.I. Lab. Technical Report* (in preparation).